

Reducing phase error in long numerical binary black hole evolutions with sixth order finite differencing

Sascha Husa, José A. González, Mark Hannam, Bernd Brügmann, Ulrich Sperhake

Friedrich Schiller University Jena, Max-Wien-Platz 1, 07743 Jena, Germany

E-mail: sascha.husa@uni-jena.de

Abstract. We describe a modification of a fourth-order accurate “moving puncture” evolution code, where by replacing spatial fourth-order accurate differencing operators in the bulk of the grid by a specific choice of sixth-order accurate stencils we gain significant improvements in accuracy. We illustrate the performance of the modified algorithm with an equal-mass simulation covering nine orbits.

PACS numbers: 04.25.Dm, 04.30.Db, 95.30.Sf

1. Introduction

In the last two years the numerical solution of the general relativistic two-body problem has made a giant leap forward with a series of breakthroughs in 2005 [1, 2, 3]. More than forty years after Hahn and Lindquist started the numerical investigation of colliding black holes [4], the field has now passed several crucial milestones toward simulating general inspiral situations, such as simulations of unequal-mass binaries and calculations of the gravitational recoil effect and the evolution of black-hole binaries with spin [5, 6, 7, 8, 9, 10]. The latter have recently lead to the spectacular finding that extremely large recoils are possible for spinning black holes [11, 12, 13].

In order to fulfill numerical relativity’s promise of providing useful information to the gravitational wave data analysis community, it is desirable to perform long numerical inspiral evolutions that allow us to cleanly match fully general relativistic and post-Newtonian waveforms with error bars, and thus produce “complete” waveforms, which contain large numbers of gravitational-wave cycles from the inspiral phase, as well as simulating the merger and ringdown phases. Such simulations will be necessary for a sufficiently dense sample of the black-hole binary configuration space.

Comparisons with post-Newtonian results have already started, and several groups have published results showing good agreement of various aspects of non-spinning simulations with post-Newtonian predictions (see e.g. [14, 15, 16, 17, 18, 19, 20]). Precise error estimates and detailed coverage of the unequal-mass and spinning cases are however missing. One serious technical problem is that performing such long evolutions with good accuracy in the phase is still computationally expensive — at least for standard “moving puncture” finite-difference codes [21, 22, 23, 16, 13, 10]. In order to overcome phase inaccuracies in long evolutions, an alternative route is

provided by spectral methods, and significant progress has been made in this direction by the Caltech-Cornell group [24, 25].

The initial data we evolve differ from the choice in [26, 25], and we cannot make a direct comparison of the accuracy of the results since puncture evolutions cannot start with excision initial data. It would certainly be worthwhile in the future to conduct direct comparisons regarding efficiency for different codes. We only quote some preliminary numbers here to illustrate the fact that while ultimately the convergence behavior of spectral codes is superior, for the purpose of “gravitational wave astrophysics”, finite difference codes may still be competitive: The performance of the Caltech/Cornell spectral code has been quoted for long time medium resolution runs as roughly 10 (for a “64³” grid configuration) and 27 (“76³” grid) CPU hours per M [26] (we will quote time and length units of the total black-hole mass M , see [16]). The highest resolution run presented here performs at 5.5 CPU hours per M on an Intel Woodcrest 2.66 GHz dual core processor. Both codes show satisfactory accuracy for long evolutions, and we certainly expect both codes to undergo further optimizations. Long evolutions that show fourth order convergence for most of the simulation have also been presented by the Goddard group [27], using impressively low spatial resolution, but no details are given on code timing.

We have previously reported accurate evolutions for approximately two orbits (initial separation of $D = 6.45M$) in [16], with an error in the merger time of 0.2 % at a computational cost of 505 CPU hours (1.44 CPU hours/ M), and we have reported a merger time error of 0.5% for $D = 8M$ simulations [28]. At larger initial separations the number of orbits is a steep function of the separation, and phase accuracy rapidly decreases. Our fourth order code would thus require resolutions which we find hard to tolerate for performing large parameter studies in the style of [6]. In the context of finite differencing it is natural to consider higher order methods. For example, it already turned out to be important to move from second-order finite differencing to fourth-order finite differencing as the feasible evolution time for puncture evolutions increased.

However, using higher order finite differencing for the types of codes used to simulate black-hole binary inspiral is not entirely trivial: for the moving-puncture method, which is currently employed in the majority of the current codes, the continuum equations become singular at the location of the “puncture”, and one might worry about the robustness of finite-difference schemes. Furthermore, current mesh refinement algorithms in the field are based on the use of buffer zones, whose number depends on the stencil width of the finite differencing scheme. In three spatial dimensions high-order finite-difference schemes with wide stencils easily lead to a drastic decrease in performance caused by the additional computational load due to the extra buffer points. This gain in accuracy pertains in particular to the phase of the evolution.

In the present paper we report on a first step to significantly improve the accuracy of current finite-difference codes to evolve black-hole binaries by using sixth order accurate finite differencing operators in the bulk of the grid. We combine the sixth order accurate derivative operators with fourth-order accurate dissipation operators [setting $r = 3$ in Eq. (3)] and Runge-Kutta time integrators, and aggressively reduce the number of AMR buffer zones compared to the number we would theoretically require for sixth order convergence. The penalty in computational cost is a rather moderate 30% compared with our fourth-order code. (We have compared the average speed over 100 M of evolution time for our largest grid configuration, which produces

a 29% increase in computational cost; in our experience this is typical for our current code).

2. Summary of the “moving-puncture method” as implemented in the BAM code

There is a large freedom in writing the Einstein equations as a system of partial differential equations, and much research has gone into finding optimal choices. In this work we employ the currently most popular choice, the BSSN system [29, 30, 31, 32], which so far is the only system for which results have been reported of long-time black-hole binary simulations that do not rely on black hole excision as has been used for example in [1, 24, 25].

We use the BSSN system together with the 1+log and gamma freezing coordinate gauges [33, 34, 31] as described in [16] (choosing in particular the parameter η in the gamma freezing shift condition as $\eta = 2$ as we have done previously). These gauge conditions allow the “punctures” to move across the grid (“moving puncture” approach [2, 3]) and allow an effective softening of the singularity in the metric associated with an internal asymptotic region [35, 36, 37], which had been prohibited by the traditional “fixed punctures” approach. The BSSN system is based on a conformal decomposition of the spatial geometry, writing the physical spatial metric as $g_{ij} = \chi^{-1} \tilde{g}_{ij}$ (following [2]). The blowup of the metric at the “punctures” is absorbed into the conformal factor χ , which vanishes at the “puncture”.

For our numerical evolutions we use the BAM [38, 39] code, which is designed to solve partial differential equations on structured meshes, in particular a coupled system of (typically hyperbolic) evolution equations and elliptic equations. The complexity of the equations is addressed by using a Mathematica package integrated into the code, which produces C-code from Mathematica expressions in tensor notation. Using such a system as we do in BAM, or as has been discussed in detail for the Cactus environment in [40] drastically simplifies the modification of complex codes for black-hole binary simulations, as was required to adapt codes from the “fixed puncture” to the “moving puncture” paradigm, or in the present case to implement the improved numerical algorithms discussed here. The structure of the BAM code has also made it straightforward to implement higher order finite differencing methods. The computational domain is decomposed into rectangular boxes, following standard domain-decomposition algorithms, and is parallelized with MPI [41]. Our mesh refinement algorithm is based on the standard Berger-Oliger algorithm, but with additional buffer zones, along the lines of [42, 43] as described in [16] and summarized in the next section. We essentially use a fixed-mesh-refinement strategy, with inner level refinement boxes following the motion of the black holes. Typically we use about 10 refinement levels (refining the grid spacing by factors of 2), roughly half of which follow the movement of the black holes.

In order to represent black holes in the initial data, we use the so-called “puncture method” [44]. For these data it is well understood how to write the constraint equations in a form suitable for numerical solution [45, 46]. Following the approach of [44] our initial data sets are chosen to be conformally flat with Bowen-York extrinsic curvature. The momentum parameter in the Bowen-York extrinsic curvature is determined from a quasi-equilibrium condition at third post-Newtonian order as described in [16]. The elliptic constraint equations are solved in BAM with the pseudo-spectral collocation code described in [47]. AMR data are then obtained by barycentric

interpolation, typically with eighth-order polynomials for both the fourth- and sixth-order finite differencing methods. The efficiency of the spectral solver is sufficient to solve the initial data problem on a single processor.

3. Sixth order finite differencing

3.1. Mesh refinement in the BAM code

Our numerical evolution algorithm is based on a method-of-lines approach using finite differencing in space and explicit fourth-order accurate Runge-Kutta time stepping (with a fixed time step). We apply sixth-order accurate polynomial interpolation in space between different refinement levels so that all spatial operations of the AMR method (i.e. restriction and prolongation) are sixth-order accurate, such that the second derivatives of interpolated values are at least fourth-order accurate. Although the time stepping used for evolution is also fourth-order accurate through the Runge-Kutta integrator, there arises the additional issue of how to provide boundary values for the intermediate time-levels of the Berger-Oliger algorithm that are not aligned in time with a coarser level (otherwise spatial interpolation can be used). Using higher than third-order interpolation has lead to spurious noise at mesh refinement boundaries as described in [16]. We therefore use third-order interpolation in time, which introduces a second-order error within the Berger-Oliger time-stepping scheme [16], which however is not noticeable in typical runs as we have checked by running with uniform (as opposed to Berger-Oliger non-uniform) time-stepping. In summary, if the outer boundary is placed sufficiently far away and if time-interpolation errors at refinement boundaries are small, then fourth-order convergence can be observed.

A relatively straightforward modification of the standard Berger-Oliger scheme is to replace the single-point refinement boundary by a buffer zone consisting of several points, e.g., [42, 43, 48]. For a sufficiently large number of buffer zones (the product of number of points in the stencil toward the mesh refinement boundary and the number of source evaluations during a full time step of the coarse grid), no time interpolation is required and excellent results have been reported for this scheme [43] (note that special methods like [49] seem to achieve similar performance). For example, our fourth-order Runge Kutta scheme requires four source evaluations, and if the lop-sided stencil with three points in one direction,

$$f'(x) = \frac{-3f_{-1} - 10f_0 + 18f_1 - 6f_2 + f_3}{12h} - \frac{1}{20}f^{(5)}(x)h^4 + O(h^5) \quad (1)$$

is used, then the numerical domain of dependence for a given point has a radius of 12 points. Here and in the following we use the notation $f_j = f(x + j\Delta x)$. Therefore, it is possible to provide 12 buffer points at the refinement boundary and to perform one RK4 time step with size three stencils that does not require any boundary updates. Only after the time step is completed, the buffer zones have to be repopulated. In the context of Berger-Oliger AMR, the buffer update is based on interpolation from the coarser levels. Since every second time step at level l coincides in time with level $l - 1$, one can provide 24 buffer points, perform two time steps, and then update the buffer by interpolation in space. With 12 buffer points, one can interpolate in time to obtain data for the buffer points at intermediate time levels.

To use fewer than 12 buffer points, we can interpolate into all buffer points before starting an RK4 update as described, and then evolve all points except the

outermost points located exactly on the boundary, which are kept fixed at their initial interpolated value. The inner points next to the boundary are updated using second-order finite differencing for the centered derivatives and shifted advection stencils for the advection derivatives. Even though for large grids the number of buffer zones becomes negligible, for the grid sizes that we use, the buffer points affect the size of the grids significantly. For example, even for our largest inner box size of 80 points in one direction, adding six points on both sides instead of 12 or 24 points leads to 92, 104, and 128 points, respectively, which corresponds to a significant saving in the number of points in 3d since $104^3/92^3 \approx 1.44$ and $128^3/92^3 \approx 2.69$. For clarity, we always quote grid sizes without buffer points, because this is the number of points owned by a particular grid. Experimentally we have found that for the fourth-order case using just six buffer points leads to very small differences compared to 12 buffer points, but even smaller buffer zones lead to noticeable differences. For simulations with fourth-order accurate derivative operators, we have therefore chosen a standard setup of RK4 with dissipation and lop-sided advection stencils, 6 buffer points, quadratic interpolation in time, and Berger-Oliger time-stepping on all but the outermost grids following [16].

As is common in numerical relativity, we use symmetric finite-difference stencils for all spatial derivatives but the advection terms associated with the shift vector, where we use lop-sided upwind stencils, see e.g. [50] for the fourth-order accurate case.

3.2. Artificial Dissipation

In finite-difference codes targeted at smooth solutions of nonlinear hyperbolic equations, it is common practice to add artificial dissipation terms to all right-hand-sides of the time evolution equations, schematically written as

$$\partial_t \mathbf{u} \rightarrow \partial_t \mathbf{u} + Q \mathbf{u}. \quad (2)$$

Such dissipation terms are very efficient at suppressing very high-frequency waves, which are not part of the physical solution. This may be necessary for numerical stability [51], but also to reduce numerical noise generated at mesh-refinement boundaries. As has become rather common in numerical relativity, we follow [51] and choose an operator (Q) of order $2r$ as

$$Q = \sigma(-h)^{2r-1}(D_+)^r(D_-)^r/2^{2r}, \quad (3)$$

for consistency with a $2r - 2$ accurate scheme, with σ a parameter regulating the strength of the dissipation. As we have done in the past with our fourth-order code we choose the factor σ as $\sigma = 0.1$ in the inner levels and $\sigma = 0.5$ in the outer levels (where the waves are extracted).

For high orders, these dissipation stencils become rather large (seven points for the fourth-order case and nine points for the sixth-order case). We therefore do not add dissipation terms where these stencils would “cross” mesh refinement boundaries. Also, adding dissipation terms with large stencils can lead to a loss of performance. We have therefore attempted to combine the use of sixth-order accurate stencils for the derivative operators with a fourth-order accurate dissipation operator and time integrator and second-order time interpolation at mesh refinement boundaries with an aggressively small number (6) of buffer points.

3.3. Sixth order accurate finite difference operators

For the sixth-order case we find that several choices for the advection term stencils yield stable evolutions, but the lop-sided upwind stencil which is closest to the symmetric case yields (probably not surprisingly) by far the best accuracy, i.e. we use

$$f'(x) = \frac{2f_{-2} - 24f_{-1} - 35f_0 + 80f_1 - 30f_2 + 8f_3 - f_4}{60h} + \frac{1}{105} \frac{d^7 f(x)}{dx^7} h^6 + O(h^7). \quad (4)$$

Alternative asymmetric choices would be

$$f'(x) = \frac{-10f_{-1} - 77f_0 + 150f_1 - 100f_2 + 50f_3 - 15f_4 + 2f_5}{60h} - \frac{1}{42} \frac{d^7 f(x)}{dx^7} h^6 + O(h^7),$$

$$f'(x) = \frac{-147f_0 + 360f_1 - 450f_2 + 400f_3 - 225f_4 + 72f_5 - 10f_6}{60h} + \frac{1}{7} \frac{d^7 f(x)}{dx^7} h^6 + O(h^7),$$

for the stencils that deviate more from the symmetric choice. We can see that the first choice has the smallest leading error term. The symmetric stencil has an even smaller error term,

$$f'(x) = \frac{-f_3 + f_{-3} - 9(f_2 - f_{-2}) + 45(f_1 - f_{-1})}{60h} - \frac{1}{140} \frac{d^7 f(x)}{dx^7} h^6 + O(h^8),$$

but does not show equally robust results, as is common for solving advection equations. For non-advection derivative terms we again use the standard symmetric stencil, similarly for second derivatives in one direction we use the symmetric stencil

$$f''(x) = \frac{-490f_0 + 270(f_1 + f_{-1}) - 27(f_2 + f_{-2}) + 2(f_3 + f_{-3})}{180h^2} - \frac{1}{560} f^{(8)}(x) h^6 + O(h^8).$$

For mixed derivatives, we use the stencils which result from a product of the symmetric sixth order accurate first derivative operators.

4. Results for long equal-mass evolutions

All runs are carried out with the symmetry $(x, y, z) \rightarrow (-x, -y, z)$ and $(x, y, z) \rightarrow (x, y, -z)$, reducing the computational cost by a factor of four. The Courant factor $C = \Delta t/h_l$ is kept constant, and is set to $C = 1/2$ for the inner grids, while for the outer grids at levels 0–4 the time step is kept constant at the value of level 3, following our previous work [16]. All runs presented here use six AMR buffer points, the same number that we have used for our fourth-order accurate code [16]. We stress that this is less than required to isolate the fine level “half” timestep from time interpolation errors at the mesh-refinement boundary, and in particular also less than required for a fully sixth-order scheme following the approach of [43]. The grid setups we have used for our simulations are displayed in table 1 (using the naming convention introduced in [16]). All the runs listed here have been performed on the Kepler cluster at the University of Jena (using Intel dual Woodcrest CPUs running at 2.66 GHz and an Infiniband interconnect), additional runs have been performed at LRZ Munich and the Doppler cluster at the University of Jena. We will denote the individual simulations by the inner-box size, i.e., 48, 56, 64, 72 or 80, as indicated in bold in table 1.

Our initial data are chosen as follows: the initial coordinate separation of the punctures is chosen as $D = 12M$, the horizon mass for each individual hole is chosen

Table 1. Grid setups used for convergence test simulations. The notation in the “Run” column is the same as we have used in [16]. The quantities h_{min} and h_{max} (rounded to three digits) denote the finest and coarsest grid spacings, and r_{max} is the location of the outer boundary (rounded to 4 digits), and all are in units of M . Also specified are the numbers of processors used, maximal memory requirement in GByte (to be precise, we quote the resident size of the program, i.e., the physical memory a task has used), and average speed in M /hour for the Kepler cluster at the University of Jena (using Intel dual Woodcrest CPUs running at 2.66 GHz). The number in bold are used to indicate individual simulations throughout this paper.

Run	h_{min}	h_{max}	r_{max}	procs.	mem. (GByte)	M/hour
$\chi_{\eta=2}[5 \times \mathbf{48} : 5 \times 96 : 6]$	1/32.0	16	776.0	8	12.2	15.6
$\chi_{\eta=2}[5 \times \mathbf{56} : 5 \times 112 : 6]$	1/37.3	96/7	774.9	12	18.2	11.6
$\chi_{\eta=2}[5 \times \mathbf{64} : 5 \times 128 : 6]$	1/42.7	12	774.0	12	22.5	8.4
$\chi_{\eta=2}[5 \times \mathbf{72} : 5 \times 144 : 6]$	1/48.0	32/3	773.3	16	31.3	3.5
$\chi_{\eta=2}[5 \times \mathbf{80} : 5 \times 160 : 6]$	1/53.3	48/5	772.8	24	45.4	4.4




Figure 1. Left panel: Coordinate tracks of the puncture location of one black hole for simulations {64, 72, 80}. Only in the last few orbits are differences between the three runs discernable. Right panel: the waveform plotted as the real part of rh_{22} , as defined in [20].

as $m_i = 0.5M$, which corresponds to puncture mass parameters of $m_p = 0.488M$. The initial momenta are obtained as $p = \pm 0.0850M$ from a 3PN-accurate quasicircularity-condition as in [16].

Our algorithm for gravitational wave extraction in terms of the Newman-Penrose scalar Ψ_4 has been described in [16]. It is useful to write the signal in terms of a time-dependent amplitude and phase as

$$\Psi_4 = A(t)e^{i\varphi(t)},$$

and define the gravitational wave frequency as $\omega = \dot{\varphi}$.

The coordinate tracks of the puncture locations are shown in figure 1 for the simulations for which we have obtained sixth-order convergence in the phase,




Figure 2. Left panel: coordinate distance of the black holes for the fourth order version of the 48-configuration and sixth-order simulations $\{48, 64, 80\}$ in the order of increasing merger time. Right panel: the gravitational wave phase for the same runs. The 72 simulation would not be distinguishable from the 80 simulation on the scale shown here.




Figure 3. Convergence test for the gravitational wave phase. Plotted are the difference between the 72 and 80 runs, and the difference between the 64 and 72 runs rescaled for sixth-order convergence. Also shown is the convergence factor divided by 6, which shows a “glitch” around the time that the phase increases very sharply, and the error estimate after performing Richardson extrapolation. The left panel shows a linear scaling, the right panel shows the same plot with a logarithmic scaling to emphasize the slow but clean exponential growth $\delta\phi = 0.0117 \exp 0.003t/M$ of the phase error at intermediate times.

$\{64, 72, 80\}$. Only during the last two orbits are differences between the three runs distinguishable. The orbital tracks show roughly 9 orbits before merger, and for the gravitational wave we obtain roughly 26 cycles before the ringdown signal becomes too noisy at $t \approx 1960M$. The right panel of figure 1 shows the real part of the $l = 2, m = 2$ mode of the rescaled strain $r h_{22}$ ($\Psi_4 = \dot{h}$, compare [20]).

Figure 2 shows the coordinate distance and gravitational wave phase for the black holes for the $\{64, 80\}$ simulations (the 72 simulation would not be distinguishable from the 80 run on the scale shown here). We find that lower resolutions merge earlier, and this is systematic for all the runs we have performed. For the $\{48, 64, 72, 80\}$ runs we obtain “merger times” of $t = (1746.8, 1790.7, 1797.5, 1800.8)M$, and the maximum of the radiation power is reached at times $t = (1818.5, 1862.2, 1869.0, 1872.4)M$. “Merger time” here is understood only as a rough indicator of the time of merger, which is used for convergence tests, and is defined as the time when the coordinate distance of the punctures drops below $0.5M$. For the $\{64, 72, 80\}$ runs the merger times show convergence at order 5.55, the radiation peak times show convergence at order 5.4.




Figure 4. The $l = 2, m = 2$ mode of the wave signal is split into the absolute value of $\Psi_{4,22}$ (left panel) and the wave frequency ω (right panel). Both panels show the simulations $\{64, 72, 80\}$, aligned in time to coincide at the peak of $|\Psi_4|$. The curves are clipped at early times, where they are very noisy due to the smallness of the signal and finite differencing error in computing the wave frequency from the phase.




Figure 5. Convergence plot for the wave amplitude $|\psi_{4,22}|$ in the $l = 2, m = 2$ mode. Both panels show the difference between the 72 and 80 runs and differences between the 64 and 72 runs rescaled for sixth-order convergence. In the left panel data at different resolutions are compared at the same coordinate time, which leads to a seeming loss of convergence near the radiation peak, which is due to the relatively large phase error. In the right panel the data are compared at the same value of the gravitational wave phase, which restores clean sixth order convergence.

Richardson extrapolation with convergence order 5.5 yields an error estimate of $\approx 4M$ for both times. Note that oscillations, which are probably mainly due to eccentricity, can clearly be seen in the black hole distance, and also in the wave amplitude shown in figure 4. A method to reduce eccentricity will be discussed in [52].

We have obtained roughly sixth order convergence for the gravitational wave phase between about $t = 1000M$ and $t = 1800M$, as shown in figure 3. At earlier times the convergence factor becomes very noisy due to the smallness of the signal. Shortly before the merger the convergence factor “glitches” to a value of roughly 7. This problem can also clearly be seen in the convergence of the radiation frequency and amplitude as shown in figures 6 and 5. This “glitch” appears when the frequency and phase increase very sharply, and small phase errors have a large effect. The logarithmic scaling version of the convergence plot in figure 3 shows a slow and rather clean exponential growth for the phase error at intermediate times, a nonlinear fit for $300M \leq t \leq 1400M$ yields $\delta\varphi = 0.0117 \exp(0.003t/M)$ for the phase error. This




Figure 6. Left panel: Reparametrisation of the black hole distance by gravitational wave phase yields clean sixth order convergence until the time of merger (which occurs roughly at the peak of the error) – the difference between the 72 and 80 runs and differences between the 64 and 72 runs rescaled for sixth-order convergence lie essentially on top of each other. Right panel: The difference in wave phase between the 4th and 6th order versions of the lowest resolution 48 configuration shows the fourth-order algorithm “falling behind”.

observation provides one way to optimize numerical methods in the inspiral phase, without evolving all the way to the merger. In the last stage of the inspiral the phase error grows very rapidly. We have noted this previously in a different context in [16], where we have compared different methods to provide quasicircular inspiral data. Small changes on the order of 1% of the initial momenta have lead to drastic changes of $\approx 40M$ in merger time. In order to clarify the convergence behaviour of our code, we have applied a new technique to check for convergence in situations where the numerical error is dominated by phase shift: We first perform a convergence analysis for the dependence of the gravitational (or orbital) phase on code time, and then perform a standard convergence test on a quantity like the puncture separation (Fig. (6)) or wave amplitude (Fig. (5)) regarding the functional dependence of this quantity on the phase.

An important question aside from convergence is how the sixth order and fourth order algorithms compare in absolute numbers. For this purpose we have re-run the 48 configuration with our standard fourth order algorithm. We find that already at this low resolution the sixth order algorithm is superior, as shown in figures 2 and 6, while at higher resolutions the larger convergence factor increases the gain in accuracy.

5. Conclusions

We have described a minimal extension of the fourth-order accurate evolution algorithm described in [16], where by replacing spatial fourth-order accurate differencing operators in the bulk of the grid by sixth-order accurate stencils, we gain drastic improvements in accuracy for the phase in long simulations of equal-mass inspiral. The crucial technical point regarding the choice of sixth-order accurate finite difference operators has been a specific choice for the advection stencil, which is used to discretize Lie derivative terms with respect to the shift vector in the Einstein equations. Using this method we have demonstrated evolutions of about nine orbits or $1800M$ with a phase error of approximately $4M$ in the merger time, requiring ≈ 11100 CPU hours on our in-house cluster. We emphasize that our code has several lower-order accurate ingredients, which however do not seem to contribute significantly to

the numerical error at the resolutions we employ.

Our emphasis here has been on boosting the current generation of “moving puncture” codes regarding their efficiency to analyse physical situations that require long evolutions, such as an accurate comparison with post-Newtonian results (see [53]), rather than on numerical analysis. Some technical questions certainly remain, such as the reduction of the numerical error at mesh refinement boundaries, the optimization for different architectures, and a rigorous mathematical analysis of numerical stability, e.g., by extending [54] to the BSSN system and the complications arising in the context of mesh-refinement. In future work we will present applications of our algorithm to other situations such as unequal masses and evolutions of spinning black holes.

Acknowledgments

This work was supported in part by DFG grant SFB/Transregio 7 “Gravitational Wave Astronomy”. We thank the DEISA Consortium (co-funded by the EU, FP6 project 508830), for support within the DEISA Extreme Computing Initiative (www.deisa.org); computations were performed at LRZ Munich and the Doppler and Kepler clusters at the Institute of Theoretical Physics of the University of Jena.

References

- [1] Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.*, 95:121101, 2005.
- [2] Manuela Campanelli, Carlos O. Lousto, Pedro Marronetti, and Yosef Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Phys. Rev. Lett.*, 96:111101, 2006.
- [3] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Gravitational wave extraction from an inspiraling configuration of merging black holes. *Phys. Rev. Lett.*, 96:111102, 2006.
- [4] Susan G. Hahn and Richard W. Lindquist. The two body problem in geometrodynamics. *Ann. Phys.*, 29:304–331, 1964.
- [5] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, J. van Meter, and M. C. Miller. Getting a kick out of numerical relativity. *Astrophys. J.*, 2007. [astro-ph/0603204](http://arxiv.org/abs/astro-ph/0603204).
- [6] J. A. González, U. Sperhake, B. Brügmann, M. Hannam, and S. Husa. The maximum kick from nonspinning black-hole binary inspiral. *Phys. Rev. Lett.*, 2006. [gr-qc/0610154](http://arxiv.org/abs/gr-qc/0610154).
- [7] M. Campanelli, C. O. Lousto, and Y. Zlochower. Gravitational radiation from spinning-black-hole binaries: The orbital hang up. *Phys. Rev. D*, 74:041501, 2006. [gr-qc/0604012](http://arxiv.org/abs/gr-qc/0604012).
- [8] M. Campanelli, C. O. Lousto, and Y. Zlochower. Spin-orbit interactions in black-hole binaries. *Phys. Rev. D*, 74:084023, 2006. [gr-qc/0608275](http://arxiv.org/abs/gr-qc/0608275).
- [9] F. Herrmann, I. Hinder, D. Shoemaker, P. Laguna, and R. Matzner. Gravitational recoil from spinning binary black hole mergers. 2007. [gr-qc/0701143](http://arxiv.org/abs/gr-qc/0701143).
- [10] Michael Koppitz, Denis Pollney, Christian Reisswig, Luciano Rezzolla, Jonathan Thornburg, Peter Diener, and Erik Schnetter. Getting a kick from equal-mass binary black hole mergers. 2007. [gr-qc/0701163](http://arxiv.org/abs/gr-qc/0701163).
- [11] J. A. Gonzalez, M. D. Hannam, U. Sperhake, B. Brügmann, and S. Husa. Supermassive kicks for spinning black holes. 2007. [gr-qc/0702052](http://arxiv.org/abs/gr-qc/0702052).
- [12] M. Campanelli, C. O. Lousto, Y. Zlochower, and D. Merritt. Large Merger Recoils and Spin Flips From Generic Black-Hole Binaries. 2007. Final version, [http://www.arxiv.org/abs/gr-qc/0701164](http://www.arxiv.org/abs/http://www.arxiv.org/abs/gr-qc/0701164).
- [13] F. Herrmann, I. Hinder, D. Shoemaker, and P. Laguna. Presented at “NRm3PN Meeting”, St. Louis, February 8-11, 2007.
- [14] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter. Binary black hole merger dynamics and waveforms. *Phys. Rev. D*, 73:104002, 2006. [gr-qc/0602026](http://arxiv.org/abs/gr-qc/0602026).
- [15] J. G. Baker, J. R. van Meter, S. T. McWilliams, J. Centrella, and B. J. Kelly. Consistency of post-newtonian waveforms with numerical relativity. 2006. [gr-qc/0612024](http://arxiv.org/abs/gr-qc/0612024).
- [16] Bernd Brügmann, José A. González, Mark Hannam, Sascha Husa, Ulrich Sperhake, and Wolfgang Tichy. Calibration of moving puncture simulations. 2006. [gr-qc/0610128](http://arxiv.org/abs/gr-qc/0610128).

- [17] Alessandra Buonanno, Gregory B. Cook, and Frans Pretorius. Inspiral, merger and ring-down of equal mass black-holes binaries. 2006. gr-qc/0610122.
- [18] Emanuele Berti et al. Inspiral, merger and ringdown of unequal mass black hole binaries: A multipolar analysis. 2007. gr-qc/0703053.
- [19] Yi Pan et al. A data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case. 2007. arXiv:0704.1964 [gr-qc].
- [20] P. Ajith et al. Phenomenological template family for black-hole coalescence waveforms. 2007.
- [21] Manuela Campanelli, C. O. Lousto, and Y. Zlochower. The last orbit of binary black holes. *Phys. Rev. D*, 73:061501(R), 2006.
- [22] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, and James van Meter. Binary black hole merger dynamics and waveforms. *Phys. Rev. D*, 73:104002, 2006.
- [23] U. Sperhake. Binary black-hole evolutions of excision and puncture data. 2006. gr-qc/0606079.
- [24] Mark A. Scheel, Harald P. Pfeiffer, Lee Lindblom, Lawrence E. Kidder, Oliver Rinne, and Saul A. Teukolsky. Solving Einstein’s equations with dual coordinate frames. *Phys. Rev. D*, 74:104006, 2006.
- [25] Harald P. Pfeiffer, Duncan Brown, Lawrence E. Kidder, Lee Lindblom, Geoffrey Lovelance, and Mark A. Scheel. Reducing orbital eccentricity in binary black hole simulations. 2007.
- [26] Mark Scheel. Binary black hole evolutions using spectral methods. Talk at the NFNR Conference, Potsdam, Germany, July 17-21 2006.
- [27] John G. Baker et al. Binary black hole late inspiral: Simulations for gravitational wave observations. 2006. gr-qc/0612117.
- [28] Bernd Brügmann, José A. González, Mark Hannam, Sascha Husa, and Ulrich Sperhake. Presented at LSC Meeting, Boston, November 4-6, 2006.
- [29] Thomas W. Baumgarte and Stuart L. Shapiro. On the numerical integration of Einstein’s field equations. *Phys. Rev. D*, 59:024007, 1999.
- [30] Masaru Shibata and Takashi Nakamura. Evolution of three-dimensional gravitational waves: Harmonic slicing case. *Phys. Rev. D*, 52:5428, 1995.
- [31] Miguel Alcubierre, Bernd Brügmann, Peter Diener, Michael Koppitz, Denis Pollney, Edward Seidel, and Ryoji Takahashi. Gauge conditions for long-term numerical black hole evolutions without excision. *Phys. Rev. D*, 67:084023, 2003.
- [32] Carsten Gundlach and Jose M. Martin-Garcia. Hyperbolicity of second-order in space systems of evolution equations. *Class. Quantum Grav.*, 23:S387–S404, 2006.
- [33] Miguel Alcubierre and Bernd Brügmann. Simple excision of a black hole in 3+1 numerical relativity. *Phys. Rev. D*, 63:104006, 2001.
- [34] John Baker, Bernd Brügmann, Manuela Campanelli, Carlos O. Lousto, and Ryoji Takahashi. Plunge waveforms from inspiralling binary black holes. *Phys. Rev. Lett.*, 87:121103, 2001.
- [35] Mark Hannam, Sascha Husa, Denis Pollney, Bernd Brügmann, and Niall Ó Murchadha. Geometry and regularity of moving punctures. 2006. gr-qc/0606099.
- [36] Mark Hannam, Sascha Husa, Niall Ó Murchadha, Bernd Brügmann, José A. González, and Ulrich Sperhake. Where do moving punctures go? *Journal of Physics: Conference series*, 2007. in press.
- [37] J. David Brown. Puncture evolution of schwarzschild black holes. 2007. arXiv:0705.1359 [gr-qc].
- [38] B. Brügmann. Binary black hole mergers in 3D numerical relativity. *Int. J. Mod. Phys. D*, 8:85, 1999.
- [39] Bernd Brügmann, Wolfgang Tichy, and Nina Jansen. Numerical simulation of orbiting black holes. *Phys. Rev. Lett.*, 92:211101, 2004.
- [40] Sascha Husa, Ian Hinder, and Christiane Lechner. Kranc: a Mathematica application to generate numerical codes for tensorial evolution equations. *Comput. Phys. Comm.*, 174:983–1004, 2006.
- [41] W. Gropp, E. Lusk, N. Doss, and A. Skjellum. A high-performance, portable implementation of the MPI message passing interface standard. *Parallel Computing*, 22(6):789–828, September 1996.
- [42] Erik Schnetter, Scott H. Hawley, and Ian Hawke. Evolutions in 3D numerical relativity using fixed mesh refinement. *Class. Quantum Grav.*, 21(6):1465–1488, 21 March 2004.
- [43] Luis Lehner, Steven L. Liebling, and Oscar Reula. AMR, stability and higher accuracy. 2005.
- [44] S. Brandt and B. Brügmann. A simple construction of initial data for multiple black holes. *Phys. Rev. Lett.*, 78(19):3606–3609, 1997.
- [45] R. Beig and N. O’Murchadha. *Class. Quantum Grav.*, 11:419, 1994.
- [46] S. Dain and H. Friedrich. Asymptotically flat initial data with prescribed regularity at infinity. *Comm. Math. Phys.*, 222:569, 2001. gr-qc/0102047.
- [47] Marcus Ansorg, Bernd Brügmann, and Wolfgang Tichy. A single-domain spectral method for

- black hole puncture data. *Phys. Rev. D*, 70:064011, 2004.
- [48] P. Csizmadia. Testing a new mesh refinement code in the evolution of a spherically symmetric Klein-Gordon field. *Int. J. Mod. Phys. D*, 15:107–119, 2006.
- [49] John G. Baker and James R. Meter. Reducing reflections from mesh refinement interfaces in numerical relativity. *Phys. Rev. D*, 72:104010, 2005.
- [50] Y. Zlochower, J. G. Baker, M. Campanelli, and C. O. Lousto. Accurate black hole evolutions by fourth-order numerical relativity. *Phys. Rev. D*, 72:024021, 2005.
- [51] Heinz Otto Kreiss and Joseph Oliger. *Methods for the approximate solution of time dependent problems*. GARP publication series No. 10, Geneva, 1973.
- [52] Sascha Husa, Mark Hannam, José A. González, Ulrich Sperhake, and Bernd Brügmann. Reducing eccentricity in black-hole binary evolutions with initial parameters from post-newtonian inspiral. In preparation.
- [53] Mark Hannam, Sascha Husa, José A. González, Ulrich Sperhake, and Bernd Brügmann. 2007. In preparation.
- [54] Gioel Calabrese, Ian Hinder, and Sascha Husa. Numerical stability for finite difference approximations of Einstein’s equations. *J. Comp. Phys.*, 218:607–634, 2005.